

WHAT IS CLAIMED:

1. A method for allowing users to define new aggregates in a database system, comprising:

receiving code that implements a class defining the structure of a user-defined aggregate and methods that can be invoked on instances of the user-defined aggregate; and

enforcing a contract against the class so that the code, when executed at runtime, satisfies requirements of the contract to ensure desired execution thereof, the contract requiring that the class comprise a first method that can be invoked to initialize the computation of an instance of the user-defined aggregate, a second method that can be invoked to accumulate a plurality of values to be aggregated, and a third method that can be invoked to compute a final result of the instance of the user-defined aggregate.

2. The method of claim 1, wherein the contract further comprises a requirement that the class specify one of a plurality of different formats for persisting instances of the user-defined aggregate in a database store.

3. The method of claim 1, wherein the contract further comprises a requirement that the class comprise a fourth method that can be invoked to merge an instance of the user-defined aggregate with another partial aggregation.

4. The method of claim 1, wherein the contract further comprises a requirement that the class comprise a public constructor having no arguments.

5. The method of claim 1, further comprising storing metadata about the user-defined aggregate for subsequent use by the database system in creating instances of the user-defined aggregate.

6. The method of claim 2, wherein the plurality of different formats for persisting instances of the user-defined aggregate comprises:

a first format in which an instance of the user-defined aggregate is automatically serialized in accordance with a native format of the database system; and

a second format in which an instance of the user-defined aggregate is serialized in a manner defined by the class.

7. The method of claim 6, wherein the plurality of different formats for persisting instances of the user-defined aggregate further comprises a third format in which an instance of the user-defined aggregate is serialized in accordance with a method provided by the MICROSOFT .NET Framework.

8. The method of claim 1, wherein the code that implements the class comprises managed code.

9. The method of claim 1, further comprising:
receiving a query that requires the evaluation of the user-defined aggregate over a group of elements;

instantiating an instance of the class that defines the user-defined aggregate;
invoking said first method to initializing the instance of the user-defined aggregate;
for each element of the group, invoking said second method on the instance of the user-defined aggregate to accumulate each element; and
invoking said fourth method to return a value of the user-defined aggregate.

10. The method of claim 9, further comprising deserializing the instance of the user-defined aggregate prior to invoking said second method for an element of the group and then serializing the instance after invoking said second method for the element.

11. The method of claim 9, further comprising:
deserializing the instance of the user-defined aggregate prior to invoking said second method for a first element of the group;
caching the deserialized instance of the user-defined aggregate;
for each element of the group, invoking said second method on the cached instance of the user-defined type; and, thereafter,
serializing the cached instance of the user-defined type.

12. The method of claim 1, wherein the class defining the user-defined aggregate further comprises one of:
an attribute that indicates whether the user-defined aggregate is invariant to duplicates;

an attribute that indicates whether the user-defined aggregate is invariant to NULL values;

an attribute that indicates whether the user-defined aggregate is invariant to order; and

an attribute that indicates whether the aggregate returns NULL if the group on which the user-defined aggregate is to be computed is empty.

13. The method of claim 12, further comprising determining a method of computation of an instance of the user-defined aggregate based at least in part upon a value of said one of said properties.

14. A database system that allows users to define new aggregates, comprising:
a runtime that provides code execution within the database system; and

a database server that receives code that implements a class defining the structure of a user-defined aggregate and methods that can be invoked on instances of the user-defined aggregate and that enforces a contract against the class so that the code, when executed at runtime, satisfies requirements of the contract to ensure desired execution thereof, the contract requiring that the class comprise a first method that can be invoked to initialize the computation of an instance of the user-defined aggregate, a second method that can be invoked to accumulate a plurality of values to be aggregated, and a third method that can be invoked to compute a final result of the instance of the user-defined aggregate.

15. The system of claim 14, wherein the contract further comprises a requirement that the class specify one of a plurality of different formats for persisting instances of the user-defined aggregate in a database store.

16. The system of claim 14, wherein the contract further comprises a requirement that the class comprise a fourth method that can be invoked to merge an instance of the user-defined aggregate with another partial aggregation.

17. The system of claim 14, wherein the contract further comprises a requirement that the class comprise a public constructor having no arguments.

18. The system of claim 14, wherein the database server stores metadata about the user-defined aggregate for subsequent use by the database system in creating instances of the user-defined aggregate.

19. The system of claim 15, wherein the plurality of different formats for persisting instances of the user-defined aggregate comprises:

a first format in which an instance of the user-defined aggregate is automatically serialized in accordance with a native format of the database system; and

a second format in which an instance of the user-defined aggregate is serialized in a manner defined by the class.

20. The system of claim 19, wherein the plurality of different formats for persisting instances of the user-defined aggregate further comprises a third format in which an instance of the user-defined aggregate is serialized in accordance with a method provided by the MICROSOFT .NET Framework.

21. The system of claim 14, wherein the code that implements the class comprises managed code.

22. The system of claim 14, wherein the database server (i) receives a query that requires the evaluation of the user-defined aggregate over a group of elements; (ii) instantiates an instance of the class that defines the user-defined aggregate; (iii) invoking said first method to initializing the instance of the user-defined aggregate; (iv) for each element of the group, invokes said second method on the instance of the user-defined aggregate to accumulate each element; and (iv) invokes said fourth method to return a value of the user-defined aggregate.

23. The system of claim 22, wherein the database server deserializes the instance of the user-defined aggregate prior to invoking said second method for an element of the group and then serializes the instance after invoking said second method for the element.

24. The system of claim 22, wherein the database server (i) deserializes the instance of the user-defined aggregate prior to invoking said second method for a first element of the group; (ii) caches the deserialized instance of the user-defined aggregate; (iii) for each element of

the group, invokes said second method on the cached instance of the user-defined type; and (iv) thereafter, serializes the cached instance of the user-defined type.

25. The system of claim 14, wherein the class defining the user-defined aggregate further comprises one of:

an attribute that indicates whether the user-defined aggregate is invariant to duplicates;

an attribute that indicates whether the user-defined aggregate is invariant to NULL values;

an attribute that indicates whether the user-defined aggregate is invariant to order; and

an attribute that indicates whether the aggregate returns NULL if the group on which the user-defined aggregate is to be computed is empty.

26. The system of claim 25, further comprising a query processor that determines a method of computation of an instance of the user-defined aggregate based at least in part upon a value of said one of said properties.

27. A computer-readable medium having program code stored thereon for allowing users to define new aggregates in a database system, the program code, when executed by a computer, causing the computer to perform the method recited in claim 1.